Journal of Statistical Software

MMMMMM YYYY, Volume VV, Issue II. *doi: 10.18637/jss.v000.i00*

arXiv:1701.05936v1 [stat.CO] 20 Jan 2017

The biglasso Package: A Memory- and Computation-Efficient Solver for Lasso Model Fitting with Big Data in R

Yaohui Zeng
University of Iowa

Patrick Breheny
University of Iowa

Abstract

Penalized regression models such as the lasso have been extensively applied to analyzing high-dimensional data sets. However, due to memory limitations, existing R packages like **glmnet** and **ncvreg** are not capable of fitting lasso-type models for ultrahigh-dimensional, multi-gigabyte data sets that are increasingly seen in many areas such as genetics, genomics, biomedical imaging, and high-frequency finance. In this research, we implement an R package called **biglasso** that tackles this challenge. **biglasso** utilizes memory-mapped files to store the massive data on the disk, only reading data into memory when necessary during model fitting, and is thus able to handle out-of-core computation seamlessly. Moreover, it's equipped with newly proposed, more efficient feature screening rules, which substantially accelerate the computation. Benchmarking experiments show that our **biglasso** package, as compared to existing popular ones like **glmnet**, is much more memory- and computation-efficient. We further analyze a 31 GB real data set on a laptop with only 16 GB RAM to demonstrate the out-of-core computation capability of **biglasso** in analyzing massive data sets that cannot be accommodated by existing R packages.

Keywords: lasso, big data, memory-mapping, out-of-core, feature screening, parallel computing, OpenMP, C++.

1. Introduction

The lasso model proposed by Tibshirani (1996) has fundamentally reshaped the landscape of high-dimensional statistical research. Since its original proposal, the lasso has attracted extensive studies with a wide range of applications to many areas, such as signal processing (Angelosante and Giannakis 2009), gene expression data analysis (Huang and Pan 2003), face recognition (Wright, Yang, Ganesh, Sastry, and Ma 2009), text mining (Li, Algarni, Albathan,

Shen, and Bijaksana 2015) and so on. The great success of the lasso has made it one of the most popular tools in statistical and machine-learning practice.

Recent years have seen the evolving era of Big Data where ultrahigh-dimensional, large-scale data sets are increasingly seen in many areas such as genetics, genomics, biomedical imaging, social median analysis, and high-frequency finance (Fan, Han, and Liu 2014). Solving the lasso efficiently for large-scale data sets remains challenging. On the other hand, Big Data itself imposes a larger challenge to the R community since R is not well-suited for analyzing large-scale data sets (Kane, Emerson, and Weston 2013). The great popularity of the lasso along with the Big Data challenges call for emerging need of efficient and scalable software tools for fitting lasso-type models.

In this project, we develop an R package, **biglasso** (Zeng and Breheny 2016), to extend lasso model fitting to Big Data in R. Specifically, sparse linear and logistic regression models with lasso and elastic net penalties are implemented. The most notable features of **biglasso** include:

- It utilizes memory-mapped files to store the massive data on the disk, only loading data into memory when necessary during model fitting. Consequently, it's able to seamlessly handle out-of-core computation.
- It is built upon pathwise coordinate descent algorithm with "warm start", "active set cycling", and "feature screening" strategies, which has been proven to be one of fastest lasso solvers. We refer the readers to Friedman, Hastie, and Tibshirani (2010) for details of the algorithm.
- We develop new, hybrid feature screening rules that outperform state-of-the-art screening rules such as the sequential strong rule (SSR) (Tibshirani, Bien, Friedman, Hastie, Simon, Taylor, and Tibshirani 2012), and the sequential EDPP rule (SEDPP) (Wang, Wonka, and Ye 2015) with additional 1.5x to 4x speedup.
- The implementation is designed to be as memory-efficient as possible by eliminating extra copies of the data created by other R packages, making **biglasso** at least 2x more memory-efficient than **glmnet**.
- The underlying computation is implemented in C++, and parallel computing with OpenMP is also supported.

The methodological innovation and well-designed implementation have made **biglasso** a much more memory- and computation-efficient and highly scalable lasso solver, as compared to existing popular R packages like **glmnet** (Friedman *et al.* 2010), **ncvreg** (Breheny and Huang 2011), and **picasso** (Ge, Li, Wang, Zhang, Liu, and Zhao 2015). More importantly, to the best of our knowledge, **biglasso** is the first R package that enables the user to fit lasso models with data sets that are larger than available RAM, thus allowing for powerful big data analysis on an ordinary laptop.

The rest of the paper is organized as follows. In Section 2, we describe the memory-mapping technique for out-of-core computation as well as our newly developed hybrid safe-strong rule for feature screening. In Section 3, we discuss some important implementation techniques as well as parallel computation that make **biglasso** memory- and computation-efficient. Section 4 presents benchmarking experiments with both simulated and real data sets. Section 5 demonstrates the out-of-core computation capability of **biglasso** through its application to a

real-world large-scale genome-wide association study. We conclude the paper with some final discussions in Section 6.

2. Method

2.1. Memory mapping

Memory mapping (Bovet and Cesati 2005) is a technique that maps a data file into the virtual memory space so that the data on the disk can be accessed as if they were in the main memory. Technically, when the program starts, the operating system (OS) will cache the data into RAM. Once the data are in RAM, the computation is at the standard in-memory speed If the program requests more data after the memory is fully occupied, which is inevitable in the data-larger-than-RAM case, OS will move currently "useless" data out of RAM to leave space for loading in new data. This is called the *page-in-page-out* procedure, which is automatically taken care of by OS.

The memory mapping technique has been commonly used in modern operating systems such as Windows and Unix-like systems due to several advantages (Wikipedia 2016). From the computational perspective, memory mapping is particularly useful for big data computation because:

(1) it provides faster file read/write than traditional I/O methods since data-copy from kernel to user buffer is not needed due to page caches;

(2) it allows random access to the data as if it were in the main memory even though it physically resides on the disk;

(3) it supports concurrent sharing in that multiple processes can access the same memory-mapped data file simultaneously, making parallel computing easy to implement in data-larger-than-RAM cases;

(4) it enables out-of-core computing thanks to the automatic page-in-page-out procedure.

We refer the readers to Rao, Prasad, and Venkateswarlu (2010), Lin, Kahng, Sabrin, Chau, Lee, and Kang (2014), and Bovet and Cesati (2005) for detailed techniques and some successful applications of memory mapping.

To take advantage of memory mapping, **biglasso** creates memory-mapped big matrix objects based upon the R package **bigmemory** (Kane *et al.* 2013), which uses the Boost C++ library and implements memory-mapped big matrix objects that can be directly used in R. Then at the C++ level, **biglasso** uses the C++ library of **bigmemory** for underlying computation and model fitting.

2.2. Efficient feature screening

Another important methodological contribution of this study is our proposed *hybrid safe-strong rule*, named SSR-BEDPP, which substantially outperforms existing state-of-the-art ones in terms of the overall computing time of obtaining the lasso solution path.

Feature screening aims to identify and discard inactive features (i.e., those with zero coefficients) from the lasso optimization. It often leads to dramatic dimension reduction and hence significant computation savings. However, these savings will be negated if the screening rule itself is too complicated to execute. Therefore, an efficient screening rule needs to be powerful enough to discard a large portion of features and also relatively simple to compute.

Existing screening rules for the lasso can be divided into two types: (1) heuristic rules, such as the sequential strong rule (SSR) (Tibshirani *et al.* 2012), and (2) safe rules, such as the basic and the sequential EDPP rules (Wang *et al.* 2015), denoted here as BEDPP and SEDPP respectively. Safe rules, unlike heuristic ones, are guaranteed to never screen a feature that should have a nonzero coefficient. Figure 1 compares the power of the three rules in discarding features. SSR, though most powerful among the three, requires a cumbersome post-convergence check to verify that it has not incorrectly discarded an active feature. The SEDPP rule is both safe and powerful, but is inherently complicated and time-consuming to evaluate. Finally, BEDPP is the least powerful, and discards virtually no features when λ is smaller than 0.45 (in this case), but is both safe and involves minimal computational burden.

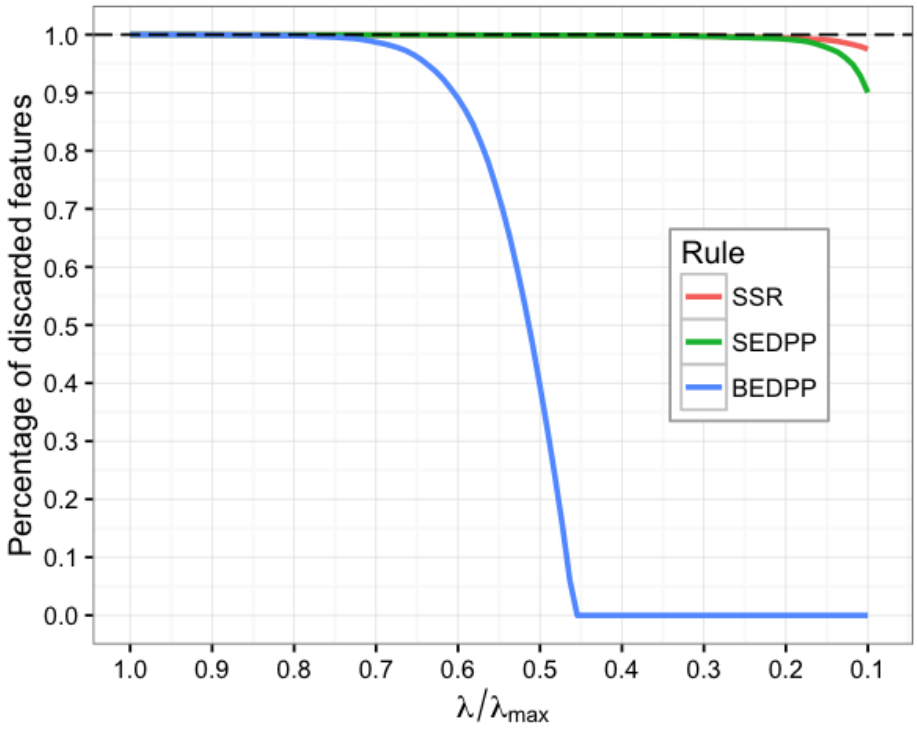

Figure 1: Percent of features discarded.

Our proposed rule, SSR-BEDPP, as its name indicates, combines SSR with the simple yet safe BEDPP rule. The rationale is to alleviate the burden of post-convergence checking for strong rules by not checking features that can be safely eliminated using BEDPP. This hybrid approach leverages the advantages of each rule, and will offer substantial gains in efficiency, especially when solving the lasso for large values of λ.

Table 1 summarizes the complexities of the four rules when applied to solving the lasso along

the entire path of $\mathcal{K}$ values of λ. SSR-BEDPP can be substantially faster than the other three rules when BEDPP is effective. Furthermore, it is important to note that SSR (with post-convergence checking) and SEDPP have to fully scan the feature matrix $\mathcal{K}$ times, while SSR-BEDPP only needs to do so for the portion of the solution path where BEDPP is ineffective. This advantage of SSR-BEDPP is particularly appealing in out-of-core computing, where fully scanning the feature matrix requires disk access and therefore becomes the computational bottleneck of the procedure.

Rule	Complexity
SSR	$O(np\mathcal{K})$
SEDPP	$O(np\mathcal{K})$
BEDPP	$O(np)$
SSR-BEDPP	$O(np + n\sum_k^{\mathcal{K}} \lvert\mathcal{S}_k\rvert))$

Table 1: Complexity of computing screening rules along the entire path of $\mathcal{K}$ values of λ. $|\mathcal{S}_k|$ denotes the cardinality of safe set $\mathcal{S}_k$; i.e., the set of features not discarded by BEDPP screening.

3. Implementation

3.1. Memory-efficient design

In penalized regression models, the feature matrix $\mathbf{X} \in \mathbb{R}^{n\times p}$ is typically standardized to ensure that the penalty is applied uniformly across features with different scales of measurement. In addition, standardization contributes to faster convergence of the optimization algorithm. In existing R packages such as **glmnet**, **ncvreg**, and **picasso**, a standardized feature matrix $\widetilde{\mathbf{X}}$ is calculated and stored, effectively doubling memory usage. This approach does not scale up well for big data.

To make the memory usage more efficient, **biglasso** doesn't store $\widetilde{\mathbf{X}}$. Instead, it saves only the means and standard deviations of the columns of $\mathbf{X}$ as two vectors, denoted as $\mathbf{c}$ and $\mathbf{s}$. Then wherever $\widetilde{x}_{ij}$ is needed, it is retrieved by "cell-wise standardization", i.e., $\widetilde{x}_{ij} = (x_{ij} - c_j)/s_j$. Additionally, the estimated coefficient matrix is sparse-coded in C++ and R to save memory space.

Simplification of computations

The "cell-wise standardization" saves a great deal of memory space, but at the expense of computational efficiency. To minimize this, **biglasso** uses a number of computational strategies to eliminate redundant calculations.

We first note that the computation related to $\widetilde{\mathbf{X}}$ during whole model fitting process are mainly of three types, and all can be simplified so that naive "cell-wise standardization" can be avoided:

(1) $\widetilde{\mathbf{x}}_j^T\widetilde{\mathbf{x}}_* = \sum_i \frac{x_{ij}-c_j}{s_j}\frac{x_{i*}-c_*}{s_*} = \frac{1}{s_j s_*}\big(\sum_i x_{ij}x_{i*} - nc_jc_*\big)$;

(2) $\widetilde{\mathbf{x}}_i^T \mathbf{y} = \sum_i \frac{x_{ij}-c_j}{s_j} y_i = \frac{1}{s_j}\big(\sum_i x_{ij} y_i - c_j \sum_i y_i\big)$;

(3) $\widetilde{\mathbf{x}}_j^T \mathbf{r} = \sum_i \frac{x_{ij}-c_j}{s_j} r_i = \frac{1}{s_j}\big(\sum_i x_{ij} r_i + c_j \sum_i r_i\big)$;

where $\widetilde{\mathbf{x}}_j$ is the jth column of $\widetilde{\mathbf{X}}$, $\widetilde{\mathbf{x}}_*$ is the column corresponding to λ_{max}, $\mathbf{y}$ is the response vector, and $\mathbf{r} \in \mathbb{R}^n$ is current residual vector.

Type (1) and (2) are used only for initial feature screening, and require only one-time execution. Type (3) occurs in both the coordinate descent algorithm and the post-convergence checking. Since the coordinate descent algorithm is fast to converge and only iterates over features in the active set $\mathcal{A}$ of nonzero coefficients, whose size is much smaller than p, the number of additional computations this introduces is small. Moreover, we pre-compute and store $\sum_i r_i$, which saves a great deal of computation during post-convergence checking since $\mathbf{r}$ does not change during this step. As a result, our implementation of cell-wise standardization requires only $O(p)$ additional operations compared storing the entire standardized matrix.

Scalable cross-validation

Cross-validation is integral to lasso modeling in practice, as it is by far the most common approach to choosing λ. It requires splitting the data matrix $\mathbf{X}$ into training and test sub matrices, and fitting the lasso model multiple times. This procedure is also memory-intensive, especially if performed in parallel.

Existing lasso-fitting R packages split $\mathbf{X}$ using the "slicing operator" directly in R (e.g., `X[1:1000, ]`). This introduces a great deal of overhead and hence is quite slow when $\mathbf{X}$ is large. What's worse, the training and test sub-matrices must be saved into memory, as well as their standardized versions, all of which result in considerable memory consumption.

biglasso, on the other hand, implements a much more stable and scalable cross-validation procedure that avoids above issues. The key design is that the main model-fitting R function allows a subset of $\mathbf{X}$, indicated by the row indices, as input. To cope with this design, all underlying C++ functions are enabled to operate on a subset of $\mathbf{X}$ given a row-index vector is provided.

Consequently, instead of creating and storing sub-matrices, only the indices of the training/test sets and the descriptor of $\mathbf{X}$ (essentially, an external pointer to $\mathbf{X}$) are needed for parallel cross validation thanks to the concurrency of memory-mapping. The net effect is that only one memory-mapped data matrix $\mathbf{X}$ is needed for K-fold parallel cross-validation, whereas other packages need (roughly) K copies of $\mathbf{X}$.

3.2. Parallel computation

Another important feature of **biglasso** is its parallel computation capability. There are two types of parallel computation implemented in **biglasso**.

At the C++ level, single model fitting (as opposed to cross validation) is parallelized with OpenMP. Though the pathwise coordinate descent algorithm is inherently sequential and thus does not lend itself to parallelization, several components of the algorithm (computing $\mathbf{c}$ and $\mathbf{s}$, matrix-vector multiplication, post-convergence checking, feature screening, etc.) do, and are parallel-enabled in **biglasso**.

By contrast, cross-validation is easily parallelized at the R level. The advantage of **biglasso**,

as compared to **ncvreg** and **glmnet**, which also support running cross-validation in parallel, is that it is much more memory- and computation-efficient by avoiding extra data copies and heavy overhead associated with copying data to parallel workers. Note that when cross-validation is run in parallel in R, parallel computing at C++ level for single model-fitting is disabled to avoid nested parallelization.

4. Numerical results

In this section, we demonstrate that our package **biglasso** (1.2-3) significantly outperforms existing popular R packages **glmnet** (2.0-5), **ncvreg** (3.7-0), and **picasso** (0.5-4). Here we focus on solving lasso-penalized linear regression over the entire path of 100 λ values which are equally spaced on the scale of λ/λ_{max} from 0.1 to 1. We set the convergence threshold to be equivalent across all four packages. All experiments are conducted with 20 replications, and the average computing times (in seconds) are reported. The benchmarking platform is a MacBook Pro with Intel Core i7 @ 2.3 GHz and 16 GB RAM.

4.1. Simulated data

Computational efficiency

We first show with simulated data that **biglasso** is more scalable in both n and p (i.e., number of observations and features). We adopt the same model in Wang *et al.* (2015) to simulate data: $\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + 0.1\boldsymbol{\epsilon}$, where $\mathbf{X}$ and $\boldsymbol{\epsilon}$ are i.i.d. sampled from $N(0, 1)$. We consider two different cases: (1) Case 1: varying p. We set $n = 1,000$ and vary p from 1,000 to 20,000. We randomly select 20 true features, and sample their coefficients from Unif[-1, 1]. After simulating $\mathbf{X}$ and $\boldsymbol{\beta}$, we then generate $\mathbf{y}$ according to the true model; (2) Case 2: varying n. We set $p = 10,000$ and vary n from 200 to 20,000. $\boldsymbol{\beta}$ and $\mathbf{y}$ are generated in the same way as in Case 1.

Figure 2 and 3, respectively, compare the mean computing time of solving the lasso over a sequence of 100 λ values by the four packages. In all the settings, **biglasso** (1 core) is uniformly 2x faster than **glmnet** and **ncvreg**, and 2.5x faster than **picasso**. Moreover, the computing time of **biglasso** can be further reduced by half via parallel-computation of 4 cores.

Memory efficiency

To demonstrate the improved memory efficiency of **biglasso** compared to existing packages, we simulate a feature matrix with dimensions $1,000 \times 100,000$. The raw data is 0.75 GB, and stored on the hard drive as an R data file and a memory-mapped file. We used `Syrupy`[1] to measure the memory used in RAM (i.e. the resident set size, RSS) every 1 second during lasso model fitting by each of the packages.

The maximum RSS during the model fitting is reported in Table 2. In the single fit case, **biglasso** consumes 0.84 GB memory in RAM, 50% of that used by **glmnet** and 22% of that used by **picasso**. Note that the memory consumed by **glmnet**, **ncvreg**, and **picasso** are respectively 2.2x, 2.1x, and 5.1x larger than the size of the raw data.

[1] https://github.com/jeetsukumaran/Syrupy

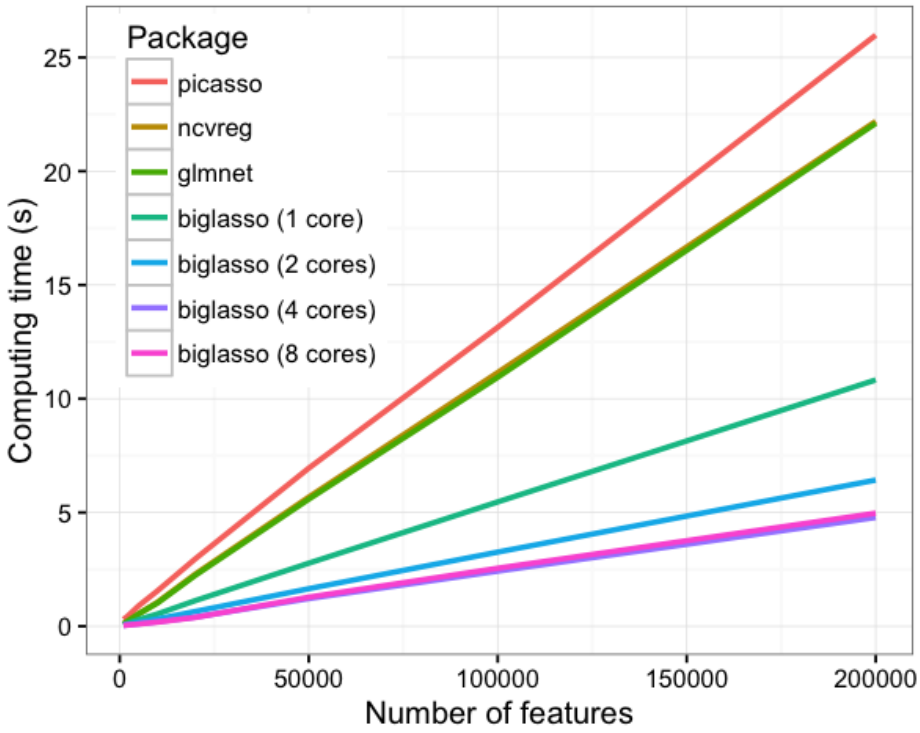

Figure 2: Mean computing time (seconds) of solving the lasso over a sequence 100 λ values as a function of p.

More strikingly, **biglasso** does not require additional memory to perform cross-validation, unlike other packages. For serial 10-fold cross-validation, **biglasso** requires just 27% of the memory used by **glmnet** and 23% of that used by **ncvreg**, making it 3.6x and 4.3x more memory-efficient than **glmnet** and **ncvreg**, respectively.

The memory savings offered by **biglasso** would be even more significant if cross-validation were conducted in parallel. However, measuring memory usage across parallel processes is not straightforward and not implemented in `Syrupy`.

Package	picasso*	ncvreg	glmnet	biglasso
Single fit	3.84	1.60	1.67	0.84
10-fold CV (1 core)	-	3.74	3.18	0.87

* Cross-validation is not implemented in **picasso**.

Table 2: The maximum RSS (in GB) for a single fit and 10 fold cross-validation (CV) with the raw data of 0.75 GB.

4.2. Real data

In this section, we compare the performance of the packages using diverse real data sets: (1)

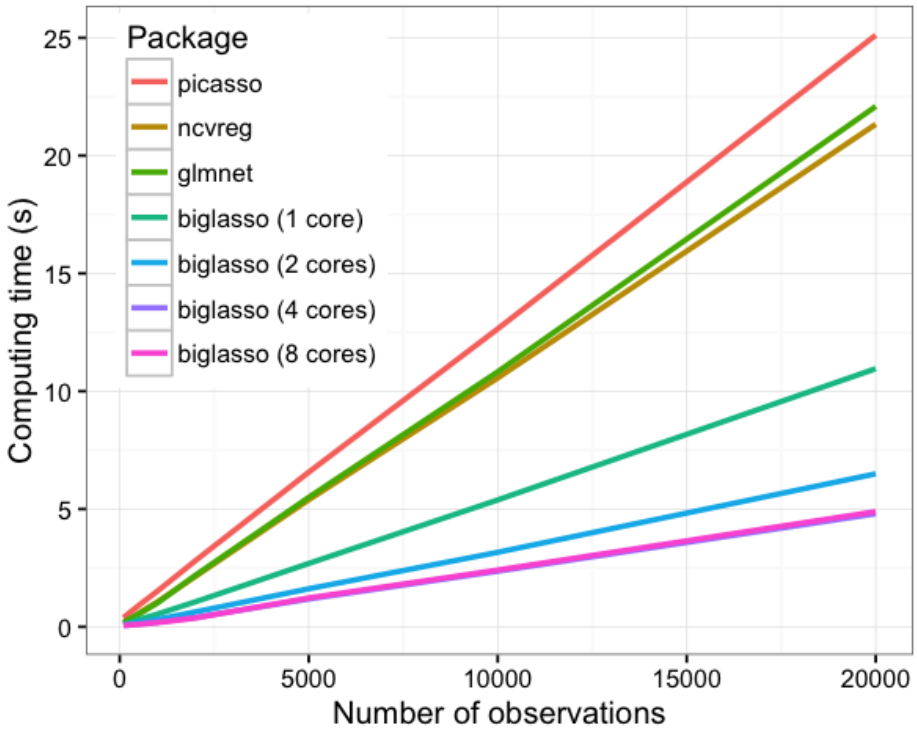

Figure 3: Mean computing time (seconds) of solving the lasso over a sequence 100 λ values as a function of n.

Breast cancer gene expression data[2] (GENE); (2) MNIST handwritten image data[3] (MNIST); (3) Cardiac fibrosis genome-wide association study data[4] (GWAS); and (4) Subset of New York Times bag-of-words data[5] (NYT). Note that for data sets MNIST and NYT, a different response vector is randomly sampled from a test set at each replication.

The size of the feature matrices and the average computing times are summarized in Table 3. Again, **biglasso** provides 2x to 3.8x speedup compared to **glmnet**, and 2x to 4.6x speedup compared to **picasso**.

5. Application: Big Data case

One of the most important features of **biglasso** is its capability of out-of-core computing. To demonstrate this, we present a real-world large-scale genome-wide association study with genetic data on 2898 infants (1911 controls and 987 premature cases). The goal is to identify genetic factors that contribute to prematurity and its complications. After preprocessing, 1,339,511 SNP (single-nucleotide polymorphism) measurements are used as features. The size of the resulting feature matrix is over 31 GB data, which is nearly 2x larger than the installed 16 GB RAM. We again solve lasso-penalized linear regression over the entire path

[2] http://myweb.uiowa.edu/pbreheny/data/bcTCGA.html
[3] http://yann.lecun.com/exdb/mnist/
[4] https://arxiv.org/abs/1607.05636
[5] https://archive.ics.uci.edu/ml/datasets/Bag+of+Words

Package	GENE $n = 536$ $p = 17,322$	MNIST $n = 784$ $p = 60,000$	GWAS $n = 313$ $p = 660,495$	NYT $n = 5,000$ $p = 55,000$
picasso	1.50 (0.01)	6.86 (0.06)	34.00 (0.47)	44.24 (0.46)
ncvreg	1.14 (0.02)	5.60 (0.06)	31.55 (0.18)	32.78 (0.10)
glmnet	1.02 (0.01)	5.63 (0.05)	23.23 (0.19)	33.38 (0.08)
biglasso	0.54 (0.01)	1.48 (0.10)	17.17 (0.11)	14.35 (1.29)

Table 3: Mean (SE) computing time (seconds) for solving the lasso along a sequence of 100 λ values.

of 100 λ values equally spaced on the scale of λ/λ_{max}. Here we consider two cases in terms of λ_{min}: (1) $\lambda_{min} = 0.1\lambda_{max}$ as before; and (2) $\lambda_{min} = 0.5\lambda_{max}$, as in practice there is typically less interest in lower values of λ for very high-dimensional data such as this.

Since other three packages cannot handle this data-larger-than-RAM case, we compare the performance of screening rules SSR and SSR-BEDPP based on **biglasso**. Table 4 summarizes the total computing time with the two rules under two cases with and without parallel computing. Clearly our proposed rule SSR-BEDPP is much more efficient in both cases. In particular, SSR-BEDPP is nearly 2.8x faster than SSR in Case 2.

Rule	$\lambda_{min}/\lambda_{max} = 0.1$		$\lambda_{min}/\lambda_{max} = 0.5$	
	1 core	4 cores	1 core	4 cores
SSR	284.56	142.55	285.61	141.28
SSR-BEDPP	189.21	93.74	102.75	51.02

Table 4: The total computing time (in minutes) of solving the lasso along a sequence of 100 λ values for 31 GB data on a laptop with 16 GB RAM.

6. Conclusion

In this research, we developed a memory- and computation-efficient R package **biglasso** to extend lasso model fitting to Big Data. Equipped with memory-mapping technique and more efficient screening rules, **biglasso** not only is 1.5x to 4x times faster than existing packages, but consumes far less memory and, critically, enables users to fit lasso models involving data sets that are too large to be loaded into memory.

Although we have discussed only lasso-penalized linear regression in this manuscript, all of the novel features we describe (memory mapping, hybrid screening rules, memory-efficient design, parallel computation) can be extended to elastic net penalties and to logistic regression, and are included in the **biglasso** package.

Since its first release in March 2016, **biglasso** has drawn increasing attention from R users. The number of its CRAN downloads is 4,016 as of January 19, 2017.

References

Angelosante D, Giannakis GB (2009). "RLS-weighted Lasso for adaptive estimation of sparse signals." In *IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 3245–3248.

Bovet DP, Cesati M (2005). *Understanding the Linux kernel*. O'Reilly.

Breheny P, Huang J (2011). "Coordinate descent algorithms for nonconvex penalized regression, with applications to biological feature selection." *Annals of Applied Statistics*, **5**(1), 232–253.

Fan J, Han F, Liu H (2014). "Challenges of big data analysis." *National Science Review*, **1**(2), 293–314.

Friedman J, Hastie T, Tibshirani R (2010). "Regularization paths for generalized linear models via coordinate descent." *Journal of Statistical Software*, **33**(1), 1.

Ge J, Li X, Wang M, Zhang T, Liu H, Zhao T (2015). ***picasso**: Pathwise Calibrated Sparse Shooting Algorithm*. R package version 0.5-4, URL https://CRAN.R-project.org/package=picasso.

Huang X, Pan W (2003). "Linear regression and two-class classification with gene expression data." *Bioinformatics*, **19**(16), 2072–2078.

Kane MJ, Emerson J, Weston S (2013). "Scalable strategies for computing with massive data." *Journal of Statistical Software*, **55**(14), 1–19.

Li Y, Algarni A, Albathan M, Shen Y, Bijaksana MA (2015). "Relevance feature discovery for text mining." *IEEE Transactions on Knowledge and Data Engineering*, **27**(6), 1656–1669.

Lin Z, Kahng M, Sabrin KM, Chau DHP, Lee H, Kang U (2014). "Mmap: Fast billion-scale graph computation on a pc via memory mapping." In *IEEE International Conference on Big Data*, pp. 159–164.

Rao ST, Prasad E, Venkateswarlu N (2010). "A critical performance study of memory mapping on multi-core processors: An experiment with k-means algorithm with large data mining data sets." *International Journal of Computers and Applications*, **1**(9), 90–98.

Tibshirani R (1996). "Regression Shrinkage and Selection via the Lasso." *Journal of the Royal Statistical Society Series B*, **58**(1), 267–288. ISSN 00359246.

Tibshirani R, Bien J, Friedman J, Hastie T, Simon N, Taylor J, Tibshirani RJ (2012). "Strong rules for discarding predictors in lasso-type problems." *Journal of the Royal Statistical Society Series B*, **74**(2), 245–266.

Wang J, Wonka P, Ye J (2015). "Lasso Screening Rules via Dual Polytope Projection." *Journal of Machine Learning Research*, **16**, 1063–1101.

Wikipedia (2016). "Memory-mapped file — Wikipedia, The Free Encyclopedia." [Online; accessed 14-September-2016], URL https://en.wikipedia.org/w/index.php?title=Memory-mapped_file&oldid=739359474.

Wright J, Yang AY, Ganesh A, Sastry SS, Ma Y (2009). "Robust face recognition via sparse representation." *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **31**(2), 210–227.

Zeng Y, Breheny P (2016). ***biglasso**: Extending Lasso Model Fitting to Big Data*. R package version 1.3-1, URL https://CRAN.R-project.org/package=biglasso.

Affiliation:

Yaohui Zeng
Department of Biostatistics
University of Iowa
N301 CPHB 145 North Riverside Drive
Iowa City, IA 52242, United States of America
Email: yaohui-zeng@uiowa.edu

Patrick Breheny
Department of Biostatistics
University of Iowa
N336 CPHB 145 North Riverside Drive
Iowa City, IA 52242, United States of America
E-mail: patrick-breheny@uiowa.edu
URL: http://myweb.uiowa.edu/pbreheny/index.html

Journal of Statistical Software http://www.jstatsoft.org/
published by the Foundation for Open Access Statistics http://www.foastat.org/
MMMMMM YYYY, Volume VV, Issue II *Submitted:* yyyy-mm-dd
doi:10.18637/jss.v000.i00 *Accepted:* yyyy-mm-dd